ESR 2: Kristoffer Andersson

Supervisors: Prof. Cornelis W. Oosterlee & Dr. Paul Koerber

Centrum Wiskunde & Informatica

Update meeting ABC-EU-XVA,
Januari 21, 2021

Overview

Overall goal and philosophy for machine learning in risk management:

- Only use machine learning when we can truly benefit from it;
- If possible, divide the problem into sub-problems, identify computational bottlenecks, and attack them with machine learning;
- Use as much mathematical structure as possible;
- Example: **Stochastic control problems**.

What have been done so far?

- Two papers related to CVA of high-dimensional derivatives with early exercise features. The first for a single derivative[1], and the second for a portfolio of derivatives[2].

What are the current topics?

- Algorithmic trading (main subject of this presentation);
- Accurate computations of derivative sensitivities in a BSDE setting (Together with a former MSc student supervised by Prof. Oosterlee and myself);
- Hedging by proxy of non-tradeable asset together with Dr. Koerber and Belfius (discussion phase).

[1] Kristoffer Andersson and Cornelis Oosterlee. "A deep learning approach for computations of exposure profiles for high-dimensional Bermudan options". In: (Mar. 2020).

[2] Kristoffer Andersson and Cornelis W Oosterlee. "Deep learning for CVA computations of large portfolios of financial derivatives". In: *arXiv preprint arXiv:2010.13843* (2020).

Table of contents

Algorithmic trading

Different types of algorithmic trading (High Frequency Trading a subset of all types):

1. **Statistical arbitrage**
   Look for anomalies in the market, *e.g.*, pairs trading;
2. **Market making**
   Provides liquidity and earns the spread (buy low - sell high and earn the difference);
3. **Optimal execution** (subject of this talk)
   How to optimally execute a large order? *e.g.*, a pension fund who wants to liquidate a large position of stocks.

   Aim to find optimal balance between:
   Sell fast $\rightarrow$ bad price low market risk
   Sell slowly $\rightarrow$ good price high market risk.

Market impact - motivation

Why not execute the trade immediately at the spot price?

Answer: Large trades affect the market, often in an unfavorable direction.

Two kinds of market impact:

- **Temporary Market Impact (TMI)**
  Liquidity dries out. Short term effect since the market re-balance quickly. Always in unfavorable direction.
- **Permanent Market Impact (PMI)**
  Signaling effect. For instance, if Warren Buffet increases his position in Apple, the stock price is likely to increase.

Active field of research, building upon[3]

[3] Robert Almgren and Neil Chriss. "Optimal execution of portfolio transactions". In: *Journal of Risk* 3 (2001), pp. 5–40.

Limit Order Book (LOB)

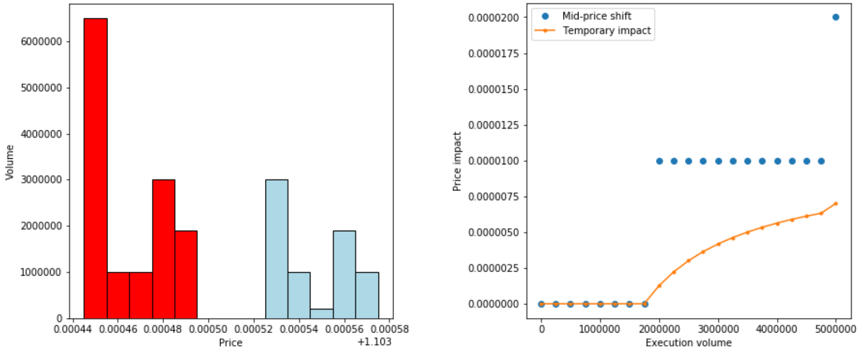

Figure: A time snapshot of **Left:** The LOB, and **Right:** Temporary Market Impact (TMI) for different trade volumes.

Temporary Market Impact (TMI)

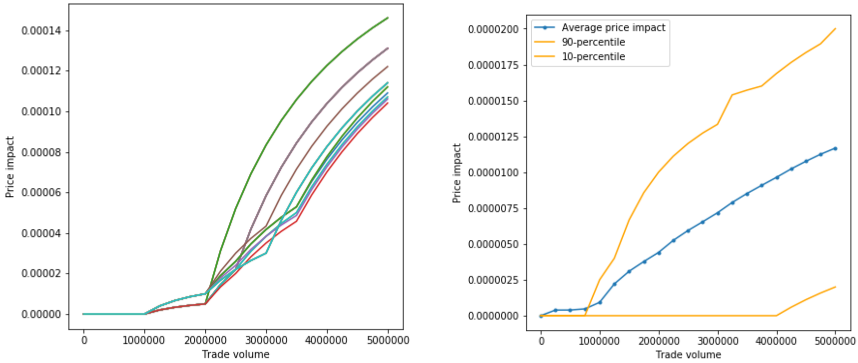

Figure: A time snapshot of **Left:** 20 different time snap shots of TMI, and **Right:** Average and 10/90-percentiles of TMI on a specific day.

Notation and setup

Assume a liquidation scheme, aiming to optimally sell Q_0 shares on the time interval $[0, T]$.

Notation:

- **Asset inventory** - $(Q_t)_{t\in[0,T]}$ (number of assets in possession);
- **Asset process** - $(S_t)_{t\in[0,T]}$ (observable asset price);
- Instantaneous **execution rate** - $(\alpha_t)_{t\in[0,T]}$;

Asset process (including PMI) given by

$$dS_t = \mu(t, S_t, \alpha_t)dt + \bar{\sigma}(t, S_t, \alpha_t)dB_t,\ \ t \in [0, T]; \quad S_0 = s_0.$$

Trading rate, time derivative of inventory

$$\frac{dQ_t}{dt} = -\alpha_t.$$

Due to TMI, we only receive

$$\tilde{S}_t = S_t - \phi(t, S_t, \alpha_t).$$

Optimization problem

Implementation shortfall given by

$$\eta_T = \overbrace{Q_0 S_0}^{\text{Initial value}} - \underbrace{\int_0^T \alpha_t \tilde{S}_t \mathrm{d}t}_{\text{Realized value}} - \underbrace{Q_T S_T}_{\text{Terminal value}} .$$

Typical objective function: **Mean-Variance**

$$J(\alpha) = \mathbb{E}[\eta_T] + \lambda \mathrm{Var}[\eta_T]$$

- Straight forward economic interpretation, close to how performance is measured in practice;
- Analytical solutions, rarely available;
- Optimal strategy time-inconsistent → No natural way to formulate the dynamic programming principle → No natural way to construct an HJB-equation and associated system of FBSDEs.

Strategy

Our approach:

1. Reformulate the problem. Want to have a state equation of the form $X = (Q, S, \eta)^T$;
   - Trading decisions based on: *i*) size of inventory, *ii*) current asset prize, and *iii*) previous performance.
2. Use the stochastic version of Pontryagin's maximum principle to formulate the adjoint equation (which gives sufficient conditions for optimality of the original problem);
   - Adjoint equation is a coupled FBSDE of McKean-Vlasov (or Mean-field) type.
3. Use modern machine learning techniques to solve the adjoint equation.
   - Make use of recently developed machine learning algorithms for FBSDEs.[4,5]

---

[4] Jiequn Han, Arnulf Jentzen, and Weinan E. "Solving high-dimensional partial differential equations using deep learning". In: *Proceedings of the National Academy of Sciences* 115.34 (2018), pp. 8505–8510.

[5] René Carmona and Mathieu Laurière. "Convergence Analysis of Machine Learning Algorithms for the Numerical Solution of Mean Field Control and Games: II–The Finite Horizon Case". In: *arXiv preprint arXiv:1908.01613* (2019).

Reformulation of state equation

Dynamics of underlying **inventory** and **asset** given by:

$$\mathrm{d}Q_t = -\alpha_t \mathrm{d}t; \quad Q_0 = q_0,$$
$$\mathrm{d}S_t = \mu(t, S_t, \alpha_t)\mathrm{d}t + \bar{\sigma}(t, S_t, \alpha_t)\mathrm{d}B_t; \quad S_0 = s_0.$$

Previous performance, described by "time integrated" version of **implementation shortfall**

$$\mathrm{d}\eta_t = -[Q_t\mu(t, S_t, \alpha_t) - \alpha_t\phi(t, S_t, \alpha_t)]\mathrm{d}t - Q_t\bar{\sigma}(t, S_t, \alpha_t)\mathrm{d}B_t; \quad \eta_0 = Q_0 S_0 = q_0 s_0.$$

Final version of state equation, $X = (Q, S, \eta)^T$, given by

$$\mathrm{d}X_t = b(t, X_t, \alpha_t)\mathrm{d}t + \sigma(t, X_t, \alpha_t)\mathrm{d}W_t, \ t \in [0, T]; \quad X_0 = (q_0, s_0, q_0 s_0)^T,$$

for appropriate b and σ and $W = (\tilde{B}, \bar{B}, B)^T$ (dummy processes $\tilde{B}, \bar{B}$).

Reformulation of objective function

Recall objective function

$$J(\alpha) = \mathbb{E}[\eta_T] + \lambda \mathrm{Var}[\eta_T].$$

Using $\mathrm{Var}[Y] = \mathbb{E}[Y^2] - (\mathbb{E}[Y])^2$ and dynamics of η gives

$$J(\alpha) = \mathbb{E}\left[\int_0^T f(t, X_t, \alpha_t)\mathrm{d}t + \lambda(\eta_t^2 - (\mathbb{E}[\eta_T])^2) \right],$$

where $f(t, X_t, \alpha_t) = -Q_t \mu(t, S_t, \alpha_t) + \alpha_t \phi(t, S_t, \alpha_t)$.

Final adjustment:

$$J(\alpha) = \mathbb{E}\left[\int_0^T f(t, X_t, \alpha_t)\mathrm{d}t + \lambda(\eta_t^2 - (\mathbb{E}[\eta_T])^2) + \underbrace{\gamma Q_T^2}_{\text{Penalizes terminal inventory}} \right].$$

Final problem formulation

Problem formulation suitable for the **Stochastic Maximum Principle (SMP)**

$$\begin{cases} \mathrm{d}X_t = b(t, X_t, \alpha_t)\mathrm{d}t + \sigma(t, X_t, \alpha_t)\mathrm{d}W_t, \ \ t \in [0, T]; \quad X_0 = x_0 \in \mathbb{R}^d, & \textbf{(State equation)} \\ J(\alpha) = \mathbb{E}\left[\int_0^T f(t, X_t, \alpha_t)\mathrm{d}t + \lambda(\eta_t^2 - (\mathbb{E}[\eta_T])^2) + \gamma Q_T^2 \right], & \textbf{(Cost functional)} \\ \text{Find } \alpha \in \mathcal{A}, \text{ such that } J(\alpha) \text{ is minimized.} & \textbf{(Problem)} \end{cases}$$

With $\mathcal{A}$ some space of admissible controls.

Hamiltonian given by

$$H(t, x, y, z, \alpha) = b(t, x, \alpha)^T y + \mathrm{Tr}(\sigma\sigma^T(t, x, \alpha)z) + f(t, x, \alpha).$$

Adjoint equation given by

$$\mathrm{d}\hat{Y}_t = -H_x'(t, \hat{X}_t, \hat{Y}_t, \hat{Z}_t, \hat{\alpha}_t)\mathrm{d}t + \hat{Z}_t\mathrm{d}W_t, \ \ t \in [0, T]; \quad \hat{Y}_T = \begin{pmatrix} 0 \\ 2\gamma\hat{Q}_T \\ 2\lambda(\hat{\eta}_T - \mathbb{E}[\hat{\eta}_T]) \end{pmatrix}.$$

Stochastic maximum principle

Theorem (necessary conditions)

Suppose some regularity conditions and that the state equation, $\hat{X} = X^{\hat{\alpha}}$ is optimally controlled by $\hat{\alpha}$ and $(\hat{Y}, \hat{Z})$ is a solution to the adjoint equation. Then for any $\bar{\alpha} \in \mathcal{A}$,

$$H'_\alpha(t, \hat{X}_t, \hat{\alpha}_t, \hat{Y}_t, \hat{Z}_t)(\hat{\alpha}_t - \bar{\alpha}_t) \geq 0, \quad \mathbb{P} - a.s. \text{ for all } t \in [0, T].$$

Theorem (Sufficient conditions)

Suppose some additional convexity assumptions. Then, if

$$H(t, \bar{X}_t, \bar{\alpha}_t, \bar{Y}_t, \bar{Z}_t) = \inf_{\alpha \in \mathcal{A}} H(t, \bar{X}_t, \alpha, \bar{Y}_t, \bar{Z}_t), \quad \mathbb{P} - a.s. \text{ for all } t \in [0, T],$$

$\bar{\alpha}$ is an optimal control and $\bar{X}$ is the optimally controlled state equation.

McKean-Vlasov FBSDE

System to solve:

$$\begin{cases} \mathrm{d}X_t = b(t, X_t, \alpha_t)\mathrm{d}t + \sigma(t, X_t, \alpha_t)\mathrm{d}W_t; \quad X_0 = x_0 \in \mathbb{R}^d, & \textbf{(Forward SDE)} \\ dY_t = -H'_x(t, X_t, Y_t, Z_t, \alpha_t)\mathrm{d}t + Z_t\mathrm{d}W_t; \quad Y_T = \begin{pmatrix} 0 \\ 2\gamma Q_T \\ 2\lambda\big(\eta_T - \mathbb{E}[\eta_T]\big) \end{pmatrix} & \textbf{(Backward SDE)} \\ H(t, X_t, \alpha_t, Y_t, Z_t) = \inf_{\alpha \in \mathcal{A}} H(t, X_t, \alpha, Y_t, Z_t) & \textbf{(Opt. cond.)}; \end{cases}$$

- Feedback-form of optimal control depend on forward SDE. With control in diffusion term - $(t, x, y, z) \mapsto \alpha$, without control in diffusion term $(t, x, y) \mapsto \alpha$;
- Coupled FBSDE (forward dynamics depend on Y and Z through α);
- McKean-Vlasov type FBSDE ($\mathbb{E}[\eta_T]$ in terminal condition of backward equation)

Algorithms - discussion

What kinds of methods are used?

1. **Analytical strategies**
   - Low flexibility on both asset dynamics and objective function (usually only minimizing implementation shortfall).
2. **PDE $\rightarrow$ Finite differences (elements)**
   - When a HJB-equation can be formulated;
   - Requires the problem to be time-consistent $\rightarrow$ strong restrictions on objective function;
   - More flexibility on asset dynamics, needs to be in low dimensions;
3. **FBSDE $\rightarrow$ Neural networks**
   - High flexibility on both asset dynamics and objective function
   - For time consistent problems reformulation through the **dynamic programming principle**. For time inconsistent problems through the **stochastic maximum principle**;
   - Mesh-free $\rightarrow$ scales good with dimensions;
   - When SMP is used, high flexibility in the action space, *e.g.*, can be a non-convex set, such as the integers.

Approximating FBSDEs with machine learning

Two most common types to approximate FBSDEs with neural networks:

- **Forward and global methods**
  - Euler-discretization in time, parametrization of Y_0 and $Z_0, Z_1, \ldots, Z_N$;
  - Loss function constructed to satisfy the terminal condition of the backward SDE.
  - Global in the sense that optimization is done only once.
- **Backward and local methods**
  - No discretization-scheme in time;
  - Approximates conditional expectations backwards in time using dynamic programming;
  - Local in the sence that optimization is done at each time point.

Extension to MV-FBSDEs (of the type considered in this presentation):

- An additional layer, since the law of the terminal state is included in the terminal condition of the backward SDE;
- Fixed point iteration until the law of the controlled terminal state converges.

Thanks for your attention!